

Introduction to Logic

[Modified Fall 2015]

I. What is logic? Claims and Arguments.

Logic is the study of correct reasoning. Here is the basic idea. Suppose that Samantha makes a *controversial claim* or has a *controversial thesis*, whether a philosophical one or some other kind. She might claim that God exists, or does not exist, that abortion is right or that abortion is wrong, that an investment is wise or that it is foolish, that the expansion of the universe is slowing down or that it is speeding up, that living things evolved from earlier organisms or not, that vaccinations are overused or that they are not. Since Samantha's claim is controversial, we ask why we should believe it. She might respond in a variety of ways. She might threaten: "Believe it or else!" Or she might look down her nose and respond disdainfully "Well, if you do not see it, you are not worth talking to." This is all well and good, but not quite what I, at least, am looking for when I ask her why I should believe what she said. What I am looking for is *evidence*. What evidence have you, Samantha, that God exists, or does not, that this investment is good or is not, that abortion is right or that it is not, that the expansion of the universe is accelerating or that it is not? For example, if Samantha tells me that abortion is morally wrong and I ask why, she might say that even the life of the unborn is important since the unborn is an innocent human being and shares all the rights of other innocent human beings. And if Samantha declares that the expansion of the universe is accelerating and I ask why we should believe it, she might talk about the red shift exhibited by distant galaxies. That is, she could respond to my questions with various kinds of evidence.

In philosophy, evidence for or against a claim is put in the form of *arguments*. In common speech, an argument is a dispute, perhaps a violent one. “Stop arguing” shouts a distraught parent after hours of disagreeable nonsense from her children. And many police reports include things like: “The argument began over a few disrespectful words and then accelerated into a general brawl.” But in logic, the word “argument” is used differently, in a *technical* sense:

An argument is a reason to believe something is true or false.

A claim is made. Someone questions the truth of the claim. Evidence is brought forward and put in the form of an argument.

In its most abstract form, an argument consists of a series of sentences. In a simple case, we have one sentence which is the *conclusion* of the argument. This is what someone wants to establish, what I earlier called a “claim” or a “thesis.” The other sentences are *premises* of the argument. They embody the evidence for the conclusion. For clarity, philosophers often put their arguments in a *standard form* which goes.

1. Premise 1
2. Premise 2
3. However many more premises there are
4. Therefore, the conclusion.

The reason for using this standard form is simple. It isolates all the premises which are supposed to be required to establish the conclusion so that there is no room for confusion about just what the evidence is supposed to be. It makes it easy to see just what the evidence is and also to see whether any evidence is missing. This, in turn, makes it easier to evaluate the argument.

Here are two common arguments dealing with the issue of abortion. At this point, I leave open whether either argument is any good. I start with a classic pro-life argument.

1. The fetus is an innocent human being from conception
2. Any innocent human being has a full right to life
3. It is always wrong to kill a being with a full right to life.
4. Abortion kills the fetus.
5. Therefore abortion is always wrong.

Next, here is a classic pro-choice argument.

1. The fetus is a tissue rather than a full-fledged human being.
2. Tissues do not have right to life and a person has the right to control what happens to tissues in her own body, even to the point of removing those tissues.
3. Therefore, a person has a right to remove the tissue that is the fetus from her body and hence to perform an abortion.

So we have competing claims, that abortion is legitimate and that abortion is not legitimate. And we have two arguments, one supporting each claim. Obviously, this is just the beginning of the discussion. We need to *evaluate* these arguments to see who has the better case. This is analogous to what happens in a court of law. Prosecutors and defense attorneys present their arguments to the jury, and the jury must decide who has *the better case*, that is, *the better arguments*. Actually, things are more complex since it is possible that one side has the better argument but that neither argument is actually very good at all! Oh well.

Something to think about: you have been given two arguments, one that supports a pro-choice position and one that supports a pro-life position on the issue of abortion. Is one of these arguments better than the other? Why?

III. What Argument is Being Given?

My discussion of arguments so far has been abstract: they consist of a series of premises and a conclusion nicely laid out. But in the real world of arguments, things are often confusing. Arguments appear in texts and in verbal discussions. I will, for

simplicity, focus on texts here. These texts can be technical works of science or philosophy, they can be newspaper editorials, magazine articles, and so on. Often the texts are disorderly. Sometimes they are downright messes. Further, many texts (perhaps most) do not provide and are not supposed to provide arguments. After all, there are many things one can do in a text. For example, a text might contain a description of events that have occurred, a prediction about events that will occur, professions of love or hate, advice, instructions, blither, and so on. So when you confront a text (and actually want to understand it) you need to ask some basic questions.

- **Is there a *claim or thesis* in this text, and if so, what is it?** As I said, many texts do not involve any claims at all. But assuming there is a claim being made in a text, sometimes it is hard to figure out just what it is. The author might not know herself, or might be a wretched writer, or might, for one reason or another, not want to state the claim explicitly.

Something to think about: Why might someone not want to state their claims clearly and explicitly? Use your imagination.

- **Is there an *argument* for this claim or thesis?** A writer might make any number of claims in a text and yet have little or no evidence for them, in which case there is not much in the way of argument. The author might not even know what an argument is or why she should give one. “That is the way I see it and if that is not good enough for you, f\$#@k off!” (Practical advice: when you come across such a text, just toss it. And if you yourself write such texts, keep them to yourself since they are not worth sharing -- unless you are in a presidential debate. Just a thought.)
- **If there is an argument, exactly what are its premises?** This can be tricky. It is common to only hint at the evidence one has and to leave out key premises. Often authors leave out premises intentionally, but just as often, authors do not realize that they need more premises to complete their arguments. Once we have all the premises, we can then put the argument into the standard form described above.

Something to think about: why might an author intentionally leave out premises, that is, part of her evidence? Can you think of examples of arguments where this has been done?

IV. What Makes an Argument “good”?

The main function of an argument is to defend a conclusion. There are two basic requirements for an argument to do this. These can be put in terms of two tests.

- A. *The Plausible Premises Test.* Arguments begin from premises. These premises embody the evidence that is meant to establish the conclusion. An argument is only as good as the premises it starts from. In logic we often assume that any premise is either true or false and say that a good argument requires true premises. But we cannot always know whether a premise is true or false. So in real life, we evaluate premises based on how *plausible* they are. If an argument begins from relatively implausible premises, we are entitled to reject the argument. The more plausible the premises, the better the argument. It is not always easy to determine how plausible a premise is.
- B. *The Follow From Test.* One hopes that one’s arguments begin from plausible premises. But just as important is whether a conclusion is actually supported by those premises, given that the premises are plausible. So, for example, it might be true that a particular pit bull is vicious. But does it follow from that fact that other pit bulls are vicious? We not only need true premises, but also ones that really do support the conclusion.

These tests are completely independent of one another An argument can do very well

on one, and terribly on the other. Consider a prosecutor in a criminal trial. He is trying to convince the jury that the defendant is guilty. To that end, he presents evidence in the form of arguments. But the prosecutors argument might pass only one of the tests. First, his argument might have very plausible premises but flunk the follow from test. For example, a prosecutor might list a number of facts and then declare they show the defendant is guilty. The defense attorney might respond that every one of those facts is correct but do little to show that her client is guilty. Second, the prosecutor might make a number of claims about the defendant which, if true, really would support his guilt, but his claims, his premises, might actually be false. The defense attorney would respond “Well, of course *if all these things were true*, my client would likely be guilty, but in fact all that stuff (the prosecutors premises) is just a pack of lies!”

So we have two tests for whether an argument is good, the plausible premise test and the follow from test. How do we evaluate an argument that does well on one test and poorly on the other? Suppose an argument gets an “A” on one test and an “F” on the other. One approach to evaluating the whole argument is to *take an average*: the average of an “A” and an “F” is a “C.” But this would be a mistake. The correct approach is the *weakest link approach*. An argument is no better than its weakest link. An argument that gets a high score on one test but flunks the other flunks on the whole. Here is an analogy. A climber examines his rope. He finds that virtually all the rope is in great shape and only one tiny bit is bad. Very bad. He declares that the rope is, on average, quite good and uses it on El Capitan. There is a word for this climber: dead. We can only hope he made his failed climb before he actually reproduced!)

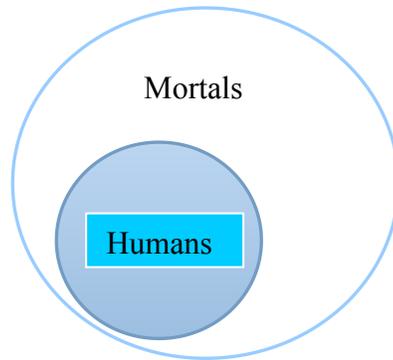
V. Induction and Deduction: two kinds of arguments and two kinds of logic.

The point behind an argument is to present premises that are plausible (The Plausible Premise Test) and which, if plausible, support the conclusion (The Follow From Test). I now focus on the follow from test. There are two radically different ways premises can support a conclusion, inductively and deductively. Roughly, premises support conclusions deductively when they do so with a kind of necessity and they support a conclusion inductively when they do so with a kind of probability. I consider these two kinds of support in turn. To this end, assume we have some true premises that are supposed to support the conclusion.

Deductive Arguments, Validity and Soundness. The idea behind a deductive argument is that if you accept the premises, *the conclusion is supposed to follow with absolute certainty.* Here is one of the simplest deductive arguments which illustrate how this works. To keep things simple, I will focus on this simple argument, but be aware that deductive arguments can be hideously complex.

1. All humans are mortal
2. Socrates is a human
3. Therefore Socrates is mortal.

This is an example from a group of arguments called “syllogisms.” It does not take much insight to see that *if those premises are true, the conclusion must be true.* If it is not perfectly obvious, the point can be made with a simple diagram. Note, this sort of diagram only works for some kinds of arguments and not for others.



The large white circle represents all mortals, including gerbils, daffodils, and humans. The smaller blue circle represents all humans. Together, the two circles represent the premise that all humans are mortal: the human circle is entirely inside the mortal circle. Obviously, if Socrates is in the human circle, he is in the mortal circle. Voila!

I have defined “deductive argument” in terms of what an argument is *supposed* to do: an argument is deductive when the premises, if true, are *supposed* to guarantee the conclusion with the kind of necessity found in the Socrates argument. But I, being rather naughty, am supposed to do lots of stuff I fail to do. And in the same way, a naughty deductive argument might fail to do what it is supposed to do. Here is an example of a naughty deductive argument very similar to the virtuous one just given. (Okay, okay, “naughty” is not a technical term from logic, but be patient.)

1. All humans are mortal
2. Socrates is a mortal
3. Therefore Socrates is human.

You should be able to see what is wrong with this argument. It is true that all humans are mortal and that Socrates is mortal, but it does not follow that he is human since there are lots of mortal things, such as gerbils and daffodils, that are not human. The truth of these premises does not guarantee the truth of the conclusion.

So why do the premises of the first argument guarantee the conclusion while the premises of the second argument do not? It turns on the *forms* of the arguments. The first has a proper form and the second does not. The first argument has the form

1. All A are B
2. x is an A
3. Therefore x is a B

The second argument has the form

1. All A are B
2. x is a B
3. Therefore x is an A.

A moments reflection on these simple forms is enough to get the point. The first form guarantees that, whatever we stick in for A and B, if the premises fit the form and are true, the conclusion must be true. The second form can provide no such guarantee.

Premises might be true and fit the form while having a false conclusion.

Consider an analogy. Think of the two forms just given as computers. Actually, this is not an analogy: they *are* computers, though not electronic ones. Just as with an electronic computer, you can feed stuff into these computers, in the form of sentences that fit the patterns given in the premises, and stuff comes out, in the form of sentences that fit the pattern in the conclusion. The first form is a good computer: feed good stuff in (that is, true premises that fit the patterns) and true stuff comes out (that is, a true conclusion that fits the pattern). The second form is a defective computer: you can feed good stuff in, but there is no guarantee that good stuff will come out. Of course, as with any defective computer, good stuff might come out, but it is hit and miss. There is no guarantee that good stuff comes out.

In sum, to be successful, any argument must pass two tests, the plausible premise test

and the follow from test. And in the context of *deductive arguments*, to be successful an argument must pass the follow from test in a very strong way: the premises must support the conclusion *with necessity* in the way illustrated by the first Socrates argument.

To bring all this together, I will introduce a few definitions that capture the most basic concepts of deductive logic. The definitions may seem a bit confusing, but you can get a feel for them by comparing them to the arguments and forms given above. The first are related to the follow from test.

- a) An argument is *valid* if and only if it has a *valid form*. Other wise it is invalid.
- b) A form is *valid* if and only if that form guarantees that any argument that fits it and has true premises also has a true conclusion. Otherwise the form is invalid.

The first deductive argument and form given above are valid. The second argument and form given above are invalid. Next, we need some terminology for deductive arguments that pass both of our tests.

- c) An argument is *sound* if and only if it is *valid* and it has *true premises*.

Otherwise, it is *unsound*.

So a sound argument is a fully successful deductive argument, one that does just what we want it to do. A deductive argument that fails to be valid, or which fails to have plausible premises, is unsound. In other words, it sucks, deductively speaking. Just to wrap things up, I will present three sucky, that is, unsound, deductive arguments. For simplicity they are variants on the arguments already given. See if you can pick out the problem with each one. There is nothing subtle here. You should get them all. Is the argument unsound because of of a false premise? Or is it unsound because of an invalid form?

1. All humans are mortal
2. Sylvester the cat is not human

3. Therefore Sylvester is not mortal.

1. All mortals are human
2. Socrates is a mortal
3. Therefore Socrates is human.

1. All dogs are reptiles
2. Socrates is a reptile
3. Therefore Socrates is dog.

Inductive Arguments, Strength and Cogency. Arguments are reasons to believe something. They are meant to start with true premises (the plausible premise test) and the premises are meant to support the conclusion (the follow from test). We have seen one way this support relation works: deductive validity. When an argument is deductively valid, the truth of the premises guarantees the truth of the conclusion, and it does this because of its form.

However, most of the arguments you come across in your lives, or come up with on your own, are not deductive. They are inductive. Obviously, a good inductive argument should start from plausible premises too. But what makes them distinctive is that the truth of the premises is supposed to render the conclusion only more or less *probable* rather than *certain*. Think of life as a series of gambles. You need to decide whether to go to college or not, whether to make this or that investment, choose this or that career, go with this or that medical diagnosis. You collect whatever evidence you can, but even if all your evidence is accurate, certainty is never possible in such matters so you can think of yourself as placing bets. Sometimes you collect evidence that might support your decision, that might make one bet look better than another. This evidence can be put in the form of an inductive argument. Such an argument might go: given all this evidence, this conclusion is more or less probable. And when you are comparing

various competing alternatives, an inductive argument might go: given all this evidence, this conclusion is more probable than that one. So, for example, you might reason “Given all this evidence, this medical diagnosis is the most plausible,” and “Given that this is the correct medical diagnosis, this treatment is most likely to be the best.”

There are many kinds of inductive argument. The most iconic is the *inductive generalization*. Here you might reason that since something has always happened a certain way in the past, it will probably (yes, that “probably” is the inductive bit) continue to happen the same way in the future. Think of inductive evidence about gravity. Roughly, objects have fallen in the past, so they will continue to fall in the future. An important variation on the inductive generalization goes like this: it has happened this way eight times out of ten in the past, so it will probably continue to happen this way eight times out of ten in the future. Another important variation involves not a *generalization* (e.g., “it will always happen this way in the future”) but a simple *prediction about the next occurrence*. So, if something has always happened a certain way in the past, we can infer that it will probably happen that way *the next time*, and if something has happened eight times out of ten in the past, we can infer that it probably has eight chances out of ten of happening this time. Here, then, are two inductive arguments which are rather important to the local ski industry.

1. It snowed in Bozeman last winter
2. It snowed in Bozeman the winter before last
3. A hundred more premises stating that it snowed in Bozeman a hundred winters in the past
4. Therefore it will probably snow in Bozeman *every winter in the future*.
1. It snowed in Bozeman last winter
2. It snowed in Bozeman the winter before last
3. A hundred more premises stating that it snowed in Bozeman a hundred winters in the past
4. Therefore it will probably snow in Bozeman *next winter*.

Something to think about: the two inductive arguments just given are similar. But one is much better than the other. Can you pick out which is the better argument and explain why it is better?

There are many other kinds of Inductive arguments. Whenever we reason that some evidence renders a conclusion more or less plausible, we are in induction-land. I will briefly explain a few more kinds of inductive arguments. There are *arguments from testimony*. These occur when we believe something because someone, some book, some document, and so on, tells us something. These arguments can look like the following.

1. My history text says that the Civil War began with the battle of Fort Sumter April 12 1861.
2. Therefore that is probably when the civil war began

1. This document, labled “Birth Certificate” and with the official seal of the State of Montana, says that my father was Fred Flintstone and my mother was Wilma Flintstone.
2. Therefore, probably, my father was Fred and my mother was Wilma

Something to think about: can you figure out which of the above inductive arguments is less convincing, and why? Which one is less convincing should be relatively easy, but it might be a bit harder to explain why it is less convincing. What exactly has gone wrong?

A huge portion of your beliefs are based on testimony. You are told by teachers, authors, friends and talking heads all sorts of stuff and you sometimes make the judgment that such testimony is more or less persuasive. Think of all the things you know, or think you know, by testimony: virtually all the history you know, except the tiny bit you yourself experienced, and almost all the science you know. In fact, you only know your

paternity by testimony! Obviously, any argument based on even the best testimony is probabilistic and the resulting arguments are inductive. “Yup, that is my dad. Huh? How do I know? My mom told me! And she would never lie about that.” (She wouldn’t? If you have any doubt that this one is probabilistic, you might want to look up the statistics, based on DNA, on incorrectly assigned paternity!)

There are also inductive arguments from *signs and symptoms*. Sometimes these are arguments from written signs, for example, highway signs. “That sign says we are entering Bozeman so I guess we probably are.” (If you doubt that this is a probabilistic argument, just reflect on the fact that Bozeman is a university town and has fraternities which may be naughty.) Other times, we speak of signs in a different sense as when we say that a particular rosy glow in Montana’s evening sky is a sign of smoke and conclude that there must be large fires somewhere. And again, when a physician collects various symptoms, and offers a probable diagnosis, she is providing an inductive argument from signs and symptoms.

When it comes time to evaluate an inductive argument, we need to ask two things analogous to the questions we asked about deductive arguments. First, are the premises really plausible? And second, if they are, do they really render the conclusion probable? Let us focus on the second question, the inductive version of the follow from test. There is a divergence here between deduction and induction. First, whether or not an argument is *deductively valid* is all or nothing. It is valid if and only if the form is valid, and otherwise it is invalid. There is no middle ground. An argument cannot be sort of valid. But when it comes to induction, things are different. True premises may supply any degree of inductive support for a conclusion. Some provide a lot of evidence, some provide a little

evidence and some provide virtually no evidence at all for the conclusion. Here are several inductive arguments for the conclusion that the defendant committed the crime. Assume the premises are true. Even assuming this, none of them provide the sort of support for the conclusion that we would want from a deductive argument. The support is only probabilistic. But some of these arguments provide better support than others.

1. This DNA specialist from the FBI forensic laboratory testified that the defendant's DNA was found on the victim.
2. Therefore, the defendant probably committed the crime.
1. This forensic hair specialist testified that the hair found on the victim is the same as the defendant's hair.
2. Therefore, the defendant probably committed the crime.
1. George claims to be an eyewitness to the crime and testified that he saw the defendant commit the crime.
2. Therefore the defendant probably committed the crime.

Something to think about: Assuming the premises are true, which of these inductive arguments is better, that is, more convincing? Why?

How can we tell how much inductive support premises provide for a conclusion? It is sometimes hard to know. In the case of deduction, there are often quite rigorous, and even quite mechanical, ways to determine whether or not an argument form is deductively valid. But there is no mechanical way to determine how much inductive support a set of premises provides for a conclusion. Often it is a matter of judgment, and even intelligent, well-informed people can disagree in their judgments. Russia sends

covert troops in to the Ukraine to support the rebels. This is evidence that there will be renewed aggression from Russia on a larger scale. The Baltic States are concerned. Finland prepares for possible war. But how strong is the evidence that further aggression will be forthcoming? Even experts differ.

I will now introduce some terminology that captures the most basic concepts in the evaluation of inductive arguments. As with deductive arguments, inductive arguments must pass the plausible premise test and the follow from test. Accordingly we have the following.

- a) An inductive argument is *strong* if and only if the truth of the premises renders the conclusion very likely. Otherwise the argument is *weak*. Obviously, there is a continuum from very strong to very weak.
- b) An inductive argument is *cogent* if and only if it is strong and its premises are in fact plausible.

So a cogent inductive argument is a fully satisfactory inductive argument. It does just what we want an inductive argument to do. It begins from plausible premises and those premises render the conclusion quite probable. .

Summary of Success Criteria for Inductive and Deductive Arguments: This has gotten complex so it will help to put it together.

- A. Successful Deductive Argument: Sound
 - a) The argument is valid, that is, the truth of the premises guarantees the truth of the conclusion because of its form.
 - b) The premises are, in fact, true.
- B. A Successful Inductive Argument: Cogent

- a) The argument is strong, that is, the truth of the premises renders the conclusion reasonably probably
- b) The premises are, in fact, true.

A final word about the terminology just provided. Words like “valid,” “sound,” “strong” and “cogent” have many meanings in ordinary English. But in logic, and therefore in philosophy, they have very specific meanings and should not be used in other ways. For example, never call a premise “strong.” That is a term reserved for certain inductive arguments. By the same token, if you want to say that an inductive or a deductive argument is fully successful, do not call it “true.” That is a term we reserve for individual statements, not arguments.

VI. Complex Arguments

In this brief introduction to logic, I presented arguments as rather simple things. A handful of premises. A straightforward conclusion. A couple of questions. Are the premises true? And if so do they support the conclusion? But in the real world arguments can be a lot more complex. I want briefly to touch on several complications.

First, it is possible to have long complex chains of reasoning where a series of premises (of which there can be many) lead to a “subconclusion.” This subconclusion can, in turn, serve as a premise (perhaps with a bunch of other new premises or subconclusions from previous arguments) to reach another subconclusion, and so on, until the final conclusion is reached. Consider the following slightly more complex chain of arguments found in discussions of the recent legalization of gay marriage. You need to ask whether each premise is plausible and whether, if plausible, they render the subconclusions, and ultimate conclusion, probable.

1. Gay marriage has just been made legal in the United States
2. This is a serious break with thousands of years of tradition.
3. Breaking with long standing traditions can seriously undermine a societies stability and its faith in its own institutions..
4. Hence (subconclusion) the stability of American society, and faith in its institutions, is likely to be seriously undermined.
5. When societies are destabilized and faith in their institutions are undermined, politicians often take extreme actions to maintain or gain power.
6. Historically, they have done this by trying to focus attention on real or imagined enemies, foreign and domestic.
7. Therefore, we can expect politicians to try to turn our attention to real or imagined enemies, foreign and domestic. (a second subconclusion)
8. The best enemies are those inside a society who are very different, or people in different societies who are already distrusted by many.
9. Therefore, we can expect many politicians to begin beating the drum of foreign war and introducing discriminatory legislation at home.

And the chain of reasoning does not stop here. We can present inductive arguments about the likely nature of these enemies, how they will be attacked and defended against, and so on.

Here is a second way an argument for a conclusion can be more complex. The argument just given, though complex, is quite linear. Though we keep adding new premises and build on previous conclusions, there is a direct line of argument from start to finish. But in many interesting cases, there are a number, perhaps a large number, of *distinct lines of argument* leading to the same conclusion. Each line of argument is independent from the others and each provides some support for the conclusion on its own. Each line might, on its own, only provide a bit of evidence for the conclusion, but taken together, the several distinct lines of argument might be strong. For example, a prosecuting attorney might stand before the jury, a scowl on his face, and maintain that there are several distinct lines of argument pointing to the guilt of the accused. No one line is sufficient to fully establish guilt, but each counts as one more “nail in the coffin.” The same thing occurs in science. For example, consider evolution. Many folk

assume that the main line of evidence supporting the existence of evolution is the fossil record. And, no doubt, some weight is put on that record. But there are a number of other independent lines of evidence as well. (If you know little biology and the following is not clear, check out most any introductory discussion of evolution on the Web.) One line turns on what are called “common structures” or “homologies.” For example, the skeletons of humans, mice, and bats are strikingly similar, despite the different ways of life of these animals and the diversity of environments in which they flourish. The same bones are found in these and many other species of animals, though a given bone might be adapted to one purpose in humans and to another in bats. This, scientists say, points toward “descent from a common ancestor with modification.” Another line or argument turns on the geographical distribution of species. There are interesting facts about how species are distributed across the planet that seem inexplicable save on evolutionary grounds. For example, there are approximately 2,000 species of beasts belonging to the genus *Drosophila* (fruit flies). About one-quarter of them live only in Hawaii. The question is why there is such a huge number of unique species of *Drosophila* on Hawaii. This odd phenomenon is easily explained on evolutionary grounds (try to guess what that explanation might be -- think islands), but seems inexplicable otherwise. Neither of these arguments is conclusive in favor of evolution, but each constitutes one more nail.

VII. Argumentative Dialogues

So far the discussion has proceeded as if arguments, whether the simple ones we started with or more complex ones, exist more or less in isolation. But interesting arguments often exist in a social environment, one that encourages a sort of dialogue

which is important in most fields, including philosophy. Such dialogues have a number of steps. Suppose that on some matter, say the existence of God, there are several theses on the table. Keeping it simple, suppose they are that God exists and that God does not exist.

1. **Step One: The Initial Positive Side of the Debate:** Each side marshals an argument, or several arguments, in favor of its preferred thesis. For example to name a few standard arguments without explaining them:
 - a) Arguments for the existence of God:
 - i. The Cosmological Argument
 - ii. The Teleological Argument
 - b) Argument against the existence of God
 - i. The Argument from Evil
2. **Step Two: The Initial Critical Side of the Debate.** Here we try to undermine our opponent's arguments. So, for example, if I do not believe in the existence of God, I will try to show what is wrong with The Cosmological and Teleological arguments. And if I am supporting the existence of God, I will try to show what is wrong with the The Argument from Evil.
3. **Step Three: The Response to Criticisms.** Here, both sides of the dispute try to undermine the criticisms offered by their opponents. I might write "I have given a cosmological argument for the existence of God. Now my opponent has responded by making points A, B and C. I, however, do not think that these points undermine my Cosmological Argument.

And no, it does not stop here. There can be responses, counter-responses,

counter-counter-responses, and so on. And there is nothing in the rule book that forbids going back and revising any of one's arguments along the line. So, for example, if I believe in the existence of God and offer a Cosmological Argument, I may come to agree with my opponent that there is something seriously wrong with my argument. Oops! Bummer. Should I leave the field of debate? Perhaps, but maybe I can modify the Cosmological Argument so as to keep it in the running. Hm. . . There are lots of ways to do it! There are many versions of the cosmological argument out there nowadays.

Why this complexity? The answer is simple: because things are complex! There are some areas in which pretty much final proofs are possible. "Here is my proof," I might say. "Not much more to be said about the matter. Good bye, Mr. and Mrs. Opponents, and turn out the lights when you leave." And I might be right to say all this. But in a great many important areas of life, things are not cut and dried. There are many competing theses being defended. There are more or less persuasive arguments for many of those theses: that is why they are all still around. There are things to be said for and against those arguments. There may currently be no winner, or if there is, that win might turn on numerous factors none of which are in themselves final. We do our best to go with the weight of the evidence knowing full well that with more thought and discussion, that weight could change. That is just the way it is. That is the way it is in many philosophical discussions, but also often in science. Even something as settled as the Big Bang might come under attack -- got that from a recent science show on TV, though I forget which one. Oh well.